

WHAT IS CLAIMED IS:

- 1 1. A method for processing data in a data store comprising:
2 obtaining a snapshot of a data store;
3 updating the snapshot with one or more first after-journal entries; and
4 after updating the snapshot with one or more first after-journal entries,
5 performing one or more subsequent updates of the snapshot with one or more second after-
6 journal entries, each subsequent update of the snapshot including:
7 storing a before-journal entry; and
8 after storing the before-journal entry, applying one of the second after-
9 journal entries to the snapshot,
10 wherein the subsequent updates of the snapshot can be undone.
- 1 2. The method of claim 1 further comprising, after performing one or
2 more subsequent updates, applying one or more before-journal entries to the snapshot,
3 wherein one or more updates of the snapshot by the second after-journal entries can be
4 undone.
- 1 3. The method of claim 2 further comprising receiving information
2 indicative of an undo request, and in response thereto performing the step of applying one or
3 more before-journal entries to the snapshot.
- 1 4. The method of claim 1 wherein the number of first after-journal entries
2 is determined based on a user-provided target time.
- 1 5. The method of claim 1 wherein the second after-journal entries are
2 applied in increasing order of time.
- 1 6. The method of claim 1 wherein the step of updating the snapshot with
2 one or more first after-journal entries includes further updating the snapshot with one or more
3 additional after-journal entries, wherein the step of further updating is performed in response
4 to receiving information indicative of a fast recovery request.
- 1 7. The method of claim 1 wherein the step of obtaining a snapshot
2 includes making a copy of the snapshot on the data store, wherein the updating steps are
3 performed on the copy of the snapshot stored on the data store.

1 8. The method of claim 1 further comprising receiving information
2 indicative of a user-specified data store, wherein the step of obtaining a snapshot includes
3 making a copy of the snapshot on the user-specified data store, wherein the updating steps are
4 performed on the copy of the snapshot stored on the user-specified data store.

1 9. A data processing device comprising:
2 a data store;
3 a controller;
4 a data storage component configured to store after-journal entries and before-
5 journal entries, and further configured to provide access to the after-journal entries and the
6 before-journal entries,
7 the controller configured to access the data store and to access the data storage
8 component,
9 the controller further configured to perform the method steps of claim 1.

1 10. A method for processing data comprising:
2 obtaining a snapshot of at least a portion of a data store;
3 applying a plurality of first after-journal entries to update the snapshot,
4 including receiving a first time indication from a user, the number of first after-journal entries
5 being based on the first time indication;
6 providing access to the snapshot so that the user can access the snapshot;
7 receiving a recovery mode indication and a second time indication from the
8 user;
9 applying a plurality of second after-journal entries to further update the
10 snapshot, the number of second after-journal entries being based on the second time
11 indication; and
12 if the recovery mode indication is indicative of an undo-able recovery mode,
13 then for each second after-journal entry, taking a before-journal entry of the snapshot before
14 applying the second after-journal entry to the snapshot.

1 11. The method of claim 10 further comprising receiving a third time
2 indication from the user and applying one or more before-journal entries to the snapshot, the
3 number of before-journal entries that are applied to the snapshot being dependent on the third
4 time indication.

1 12. A data processing system comprising:
2 a host component comprising at least one host processing unit;
3 a storage component comprising at least one storage control unit;
4 first program control means contained in the host component for controlling
5 operation of the host processing unit; and
6 second program control means contained in the storage component for
7 controlling operation of the storage control unit,
8 the first program control means and the second program control means further
9 for operating, respectively, the host processing unit and the storage control unit to perform
10 the method steps of claim 10.

1 13. The data processing system of claim 12 wherein the first program
2 control means comprises first program code and the second program control means comprises
3 second program code.

1 14. A method for processing data on a data store comprising:
2 receiving input from a user indicative of a first data volume;
3 receiving input from the user indicative of a second data volume;
4 obtaining a snapshot of at least a portion of the first data volume;
5 storing the snapshot on the second data volume;
6 a first step of updating the snapshot with a plurality of first after-journal
7 entries;
8 providing user-access to the second data volume;
9 receiving a first indication from the user, wherein if the first indication is
10 indicative of a fast recovery operation, then repeating the first step of updating the snapshot
11 with a plurality of second after-journal entries; and
12 subsequent to the first step of updating, a second step of updating the snapshot
13 with a plurality of third after-journal entries, including for each third after-journal entry
14 taking a before-journal entry of the snapshot prior to updating the snapshot with the third
15 after-journal entry,
16 the first, second, and third after-journal entries being representative of write
17 operations previously performed on the first data volume.

1 15. The method of claim 14 further comprising receiving input from the
2 user indicative of a target time wherein the number of first after-journal entries is based on
3 the target time.

1 16. The method of claim 15 further comprising receiving input from the
2 user indicative of a refined target time wherein the number of second after-journal entries is
3 based on the refined target time.

1 17. The method of claim 15 further comprising receiving input from the
2 user indicative of a refined target time wherein the number of third after-journal entries is
3 based on the refined target time.

1 18. The method of claim 14 further comprising applying one or more
2 before-journal entries to the snapshot to undo snapshot updates produced by the application
3 of one or more of the third after-journal entries.

1 19. The method of claim 14 further comprising receiving a second
2 indication from the user and in response thereto, applying one or more before-journal entries
3 to the snapshot to undo snapshot updates produced by the application of one or more of the
4 third after-journal entries.

1 20. The method of claim 19 further comprising receiving input from the
2 user indicative of a time, wherein the number of before-journal entries is based on the time.

1 21. The method of claim 19 wherein the one or more before-journal entries
2 are applied sequentially beginning with the most recent before-journal entry.

1 22. The method of claim 14 wherein the first data volume and the second
2 data volume refer to the same data volume, wherein the snapshot represents a data state of at
3 least a portion of the first data volume at a first point in time.

1 23. The method of claim 14 wherein the first data volume is a production
2 volume and the second data volume refers to a data volume different from the production
3 volume.